

APPLICATION FOR UNITED STATES PATENT

For

**Method and Apparatus For Efficient SPVC Destination Endpoint Address
Change**

Inventors:

Krishna Sundaresan

Mahesh Chellappa

Daniel Cauchy

Prepared by:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP
12400 Wilshire Boulevard
Los Angeles, CA 90025-1030
(408) 720-8300

Attorney's Docket No.: 081862.P251

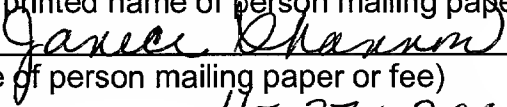
"Express Mail" mailing label number: EL351963063US

Date of Deposit November 27, 2001

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Commissioner for Patents, Arlington, VA 22202.

Janece Shannon

(Typed or printed name of person mailing paper or fee)


(Signature of person mailing paper or fee)

11-27-2001
(Date signed)

10/22/01 08:49:55

Method and Apparatus For Efficient SPVC Destination Endpoint Address Change

FIELD OF THE INVENTION

[0001] The field of invention relates to networking, generally; and, more specifically, to a method and apparatus for efficient SPVC destination endpoint address change.

BACKGROUND

[0002] An exemplary Private Network Node Interface (PNNI) Asynchronous Transfer Mode (ATM) network 101 is shown in **Figure 1**. ATM is a networking technology that transports information with “cells” of data. As such, if a significantly sized body of information (e.g., a document or file) is to be transported across an ATM network, the body of information is effectively “broken down” into a plurality of cells. The plurality of cells are then individually sent across the network and reassembled at the receiving end in order to reconstruct the original body of information.

[0003] The term “connection” or “circuit” is often used to describe a pre-defined path through a network. Typically, when a body of information is to be transported over a network, a connection is setup beforehand that establishes (in some manner and to some extent) the path that the cells will take. Various types of connections may be used within an ATM network 101. These include: 1) permanent virtual circuits (PVCs); 2) switched virtual circuits (SVCs); and 3) soft permanent virtual circuits (SPVCs).

[0004] In the case of PVCs, a quasi-permanent connection is established (e.g., a connection that lasts for days, weeks, months, etc.). PVCs are often used in situations where a large corporate user desires to permanently clear a guaranteed pipe through the network 100 from one large office to another large office. For example, if node 105₁ corresponds to the Customer Premise Equipment (CPE) of a first corporate office and node 105₂ corresponds to the CPE of a second corporate office, a PVC may be established that couples nodes 102₁, 102₄, 102₇ and network lines 103₃, 103₁₁ together (in order to form an end-to-end path through the network 100 between CPEs 105₁ and 105₂).

[0005] Generally, the amount of traffic (e.g., as between two large corporate offices) and the extent of the usage (e.g., every business day for the foreseeable future) justifies the costs associated with dedicating, in a quasi-permanent fashion, a fixed amount of the network's resources to one particular pathway. Typically, a PVC is manually configured by a network manager from a network management control station 104. As such, commands are issued from the network control station 104 to the various nodes in the network 101 that "make up" the PVC (so that the lookup tables, etc. within these nodes can be properly updated).

[0006] Another characteristic of a PVC is that a PVC user simply directs traffic into the network 101 (e.g., from node 105₁) with little or no formal request for transportation services from the network 101. For example, typically, a PVC user at node 105₁ will send ATM cells having the PVC's VPI/VCI across the ATM User Network Interface (UNI) at link 103₁. Based upon the VPI/VCI information, node

102₁ (e.g., as well as subsequent nodes along the PVC path) will be able to properly switch the cells onto a link that corresponds to the PVC path. Thus, because the connection is quasi-permanent and has already been established, there is little or no procedural overhead associated with connection setup (such as a SETUP request message and the like). The user is provided an appropriate VPI/VCI well beforehand (e.g., shortly after PVC setup) which is invoked each time thereafter by the user when the services of the PVC are desired.

[0007] SVCs, on the other hand, are established on a temporary basis rather than a quasi-permanent basis. SVCs efficiently utilize the resources of a network if the network has to support a large number of different connection paths over a fairly brief period of time (e.g., seconds, minutes, hours). In contrast to PVCs, SVCs are usually established on a “call-by-call” basis and therefore have: 1) some form of formal user request to the network 101 for transportation services; and, 2) a connection “setup” procedure that follows the request for transportation services and a connection “teardown” procedure that follows the successful performance of the requested transportation services.

[0008] The connection setup/teardown procedures may be viewed as the “automatic” configuration of a connection within the network rather than manual configuration from a network management control station 104. PNNI is a routing and signaling protocol that determines and establishes connection paths . The PNNI routing protocol is executed on the source endpoint (e.g., source endpoint 102₁ for connections initiated from originating node 105₁), and is often referred to

as a “source” routing protocol. An example of PNNI’s routing and signaling techniques are provided immediately below.

[0009] If node 105₁ (the “originating” node) desires to send information to node 105₂ (the “target” node), the originating node 105₁ will effectively request the network 101 for a connection to be established between nodes 105₁ and node 105₂. Typically, this request takes the form of a SETUP message that is passed over the ATM UNI at link 103₁. The access node 102₁ (which may be referred to as the source endpoint node) receives the SETUP message and determines an appropriate path for the connection through the network via the PNNI routing protocol.

[0010] The SETUP message then traverses the network 101 to the destination endpoint node 102₇. When the SETUP message is received at the destination endpoint node 102₇, a CONNECT message is issued from the destination endpoint node 102₇ to the source endpoint node 102₁. The CONNECT message “bounces”, node-by-node, along the connection path to the source endpoint node 102₁. Each node that receives the CONNECT message updates its lookup table (or other routing/switching platform) with an appropriate reference to the connection being established. When the source endpoint node 102₁ receives the CONNECT message, the VPI/VCI for the connection is passed to the user at the ATM UNI (along link 103₁), the connection is established, and transportation services may commence. After the transportation services are complete, the connection is torndown in a manner similar to that in which it was established.

[0011] An SPVC is often viewed as a blending of an SVC and a PVC. SPVCs are often used to provide guaranteed bandwidth to a particular user (such that the user enjoys service as if a permanent pipe has been established through the network 101) while, simultaneously, the network 101 is allowed to flexibly adapt to different connection paths over brief periods of time (by establishing each SPVC communication with connection setup and teardown procedures). In order to implement an SPVC service, the endpoint nodes of the ATM network 101 (e.g., source node 102₁ and destination node 102₇) are configured to behave like PVC nodes with respect to the user (e.g., along the ATM UNI at link 103₁) while behaving like SVC nodes within the ATM network 101 itself.

[0012] With an SPVC, the source and destination endpoint nodes 102₁ and 102₇ are usually manually configured by the network management station 104 to provide a PVC interface to the users at node 105₁ (and at node 105₂). That is, for example, a quasi permanent VPI/VCI is provided to the user that is to be invoked each time the services of the SPVC are desired. Upon the receipt of ATM cells having this VPI/VCI information, however, the endpoint source node 102₁ triggers the release of a SETUP message which traverses the network 101 to destination endpoint node 102₇. A CONNECT message is returned to the endpoint source node 102₁, and the SPVC is established.

Figures

[0013] The present invention is illustrated by way of example, and not limitation, in the Figures of the accompanying drawings in which.

[0014] **Figure 1** shows an embodiment of a PNNI ATM network.

[0015] **Figure 2** shows an embodiment of a PNNI Topology State Packet (PTSP).

[0016] **Figure 3** shows an embodiment of a PNNI Topology State Element (PTSE) that may be embedded within the PTSP of **Figure 2**.

[0017] **Figure 4** shows an embodiment of a System Capabilities Information Group (SIG) field that may be embedded within the PTSE of **Figure 3**.

[0018] **Figure 5** shows an embodiment of a methodology that may be executed by an SPVC source endpoint in order to reconfigure itself so that a change in the SPVC destination endpoint can be recognized.

[0019] **Figure 6** shows an embodiment of a node.

Description

[0020] A problem with SPVC connections is the inefficiencies associated with changing the address of an endpoint destination node. That is, each node 102₁ through 102₇ is referenced according to its own unique address. Examples include the NSAP addressing format or the E.164 addressing format. If the address of the endpoint destination node changes, a change should be made to each source node that handles a PVC or SPVC that is directed to the particular endpoint destination node whose address is being changed.

[0021] For example, using the SPVC example referred to above as a basis for discussion, if destination endpoint destination node 102₇ is to undergo a change in address value then source endpoint node 102₁ should reflect this change so that SETUP messages for subsequent SPVC connections will be properly directed to destination endpoint node 102₇. As discussed, the configuration of an SPVC endpoint node is typically performed via manual efforts that are exerted from the network management station 104.

[0022] As such, the SPVC information of source endpoint node 102₁ will be manually reconfigured to reflect the address change of destination endpoint 102₇. Furthermore, to the extent that node 102₇ acts as a destination endpoint node for other SPVCs within network 101, the corresponding source endpoint nodes for each of these SPVCs should be similarly reconfigured. For example, if nodes 102₂, 102₃, 102₅, and 102₆ each behave as a source endpoint node for an SPVC that is directed to node 102₇, each of these nodes 102₂, 102₃, 102₅, and 102₆ will

also be manually reconfigured to reflect a change in the destination endpoint node 102₇.

[0023] In complex networks where a single node can act as the destination endpoint for hundreds or thousands of SPVCs, an extensive manual effort may be required to reconfigure the source endpoint of each of these SPVCs. The result is high network maintenance and management costs. The inefficiencies associated with the changing of a destination endpoint address can be improved, however, by building a mechanism into the network 101 that automatically reconfigures each SPVC source endpoint that is affected by a change in an SPVC destination endpoint node address change.

[0024] Because the reconfiguration of each affected SPVC source endpoint node is automatic, the reconfiguration can be successfully completed in the absence of manual efforts that are directed from the network management control station 104. As such, an improvement in network management efficiency is realized. Automatic reconfiguration may be accomplished via the use of PNNI Topology State Elements (PTSEs) which are described in more detail below.

[0025] As discussed, PNNI is a routing and signaling protocol that is executed on each node in the network 101. As part of the PNNI scheme, each node is designed to "broadcast" information pertaining to its understanding of itself and/or the network in which it resides. These broadcasts may occur at specific time intervals and/or upon the occurrence of certain special events.

[0026] For example, if a node 102₅ observes that networking link 103₁₀ is not working, the node 102₅ will broadcast this event to its neighboring nodes 102₂,

102₇. Upon the reception of this information, the neighboring nodes 102₂, 102₇ will “update” their internal understandings of the network (to reflect this event) as well as rebroadcast this event to their neighboring nodes so that they may update their internal understandings as well. The information is continually rebroadcast as appropriate so that the affected nodes can update their understandings of the network and behave accordingly.

[0027] Thus, in a sense, the occurrence of the event ripples through the network so that its constituent nodes can cohesively route information around the downed link 103₁₀ in response. In other cases, typically, the network’s nodes 102₁ through 102₇ are also configured to broadcast current status information as well as special events. Thus, on a broader scale, the nodes of the network may be said to communicate procedural (e.g., “control”) information with one another as well as the substantive information associated with user traffic.

[0028] This control information is often organized into one or more PNNI Topology State Elements (hereinafter, referred to as PTSEs) that are embedded into a PNNI Topology State Packet (hereinafter, referred to as a PTSP). A PTSP is a packet that acts as the broadcast mechanism while a PTSE acts as a component of the PTSP’s payload. Thus, for example, if a node has information to broadcast it issues a PTSP that carries one or more PTSEs that each have the information to be communicated. An embodiment 200 of a PTSP is shown in **Figure 2** and an embodiment 301 of a PTSE is shown in **Figure 3**.

[0029] Referring to **Figure 2**, a PTSP may be viewed as having a header field 206 and a PTSE field 201. The header field 206 has various header information

(e.g., checksum info, lifetime, etc.) as well as the identification of the node that is issuing the PTSP (which is located within the originating node ID field 203), the peer group within which the originating node resides (which is located within the Peer Group ID field 204). PNNI Peer groups are discussed in more detail toward the end of this description.

[0030] The PTSE field 201 includes one or more PTSEs 201₁ through 201_x. An embodiment 301 of a PTSE is shown in **Figure 3**. That is, for example, the PTSE embodiment 301 of **Figure 3** may be viewed as corresponding to the PTSE 201₁ of **Figure 2**. Referring to **Figure 3**, note that a PTSE may also be viewed as having a header field 302 and a payload field 203. The header field 302 includes various header information such as a type field 306 that identifies the data structure 301 as a PTSE, a length field 307 that identifies the length of the PTSE, a reserved field 309 for potential future uses and a checksum field 312.

[0031] The PTSE header field 302 also includes a identifier field 310 that identifies the type of PSTE that PTSE 301 corresponds to. That is, PNNI employs a characterization scheme so that specific types of information can be binned together or recognized within a common PTSE format. The various PTSE types include (among possible others): 1) Horizontal Link; 2) Uplink; 3) External Address; 4) Internal Address; 5) Nodal Parameters (Complex Node); and 6) Nodal. Those of ordinary skill can identify the purpose and/or use of each PTSE type.

[0032] However, it is noteworthy to point out that the “Nodal” PTSE type is typically used to broadcast status information about the node that originates the PTSE. As such, it is an appropriate PTSE type for broadcasting a change in an SPVC destination endpoint address. Specifically, in one embodiment, the PNNI scheme is extended such that any node which experiences both an address change and acts as a destination endpoint for one or more SPVCs is configured to issue a PTSP having a “Nodal” PTSE that includes information which indicates that an address change is at hand.

[0033] Referring to the PTSE embodiment 301 of **Figure 3**, note that the payload field 303 may be viewed as being partitioned into an “industry standard” field 304 and a System Capabilities Information Group (SIG) field 305. The industry standard field 304 is used to carry specific information according to a specific format that has been articulated by the PNNI standard. The SIG field 305, by contrast, is used for developers of PNNI compliant networking gear that seek to include special features beyond those recognized or articulated by the PNNI standard.

[0034] Through the use of the SIG field 305, two nodes from the same manufacturer can communicate information with one other that is not specifically provided for by the PNNI standard; while, at the same time, operate in compliance with the PNNI standard. That is, those nodes that can understand and use the contents of the SIG field 305 may do so while those that do not understand the SIG field 305 contents may simply ignore its information (as well

as forward the PTSE having the SIG field to another node via a rebroadcast effort).

[0035] Figure 4 shows an embodiment 405 of a SIG field. That is, the SIG field 405 of Figure 4 may be viewed as an embodiment of the SIG field 305 of Figure 3 that can be used to express the address change of an SPVC endpoint. The SIG field embodiment 405 of Figure 4 can also be viewed as having a header field component 401 and a payload field component 402.

[0036] The header field component 401 includes various header information such as a type field 406 (that indicates the data structure 405 is a SIG field), a length field 407 that describes its length and a an Organization Unique Identifier (OUI) field 408 that is typically used to recognize the manufacturer of the node that issued the SIG information (i.e., is a “vendor-specific” label). As a SIG field is typically used by the nodes of a common manufacturer to support functional improvements (beyond the PNNI standard) that are unique to the products of the manufacturer, the OUI field 408 is often used by a node to decide whether or not to ignore a received SIG field. That is, if the vendor specific label of the OUI field 408 “matches” the vendor of the node that receives the SIG information, the SIG information will be “looked into”; otherwise, the SIG information will be discarded.

[0037] Within the payload 402 of the SIG field 405, the ID # field 403 identifies the particular type of information being delivered by the SIG 405. This allows a node that supports vendor-specific functionality to understand the specific type of information enclosed in the payload 402. As such, in an embodiment, a specific binary number is used to identify that the SIG field 405 includes information

related to the address change of an SPVC endpoint destination node. In the particular embodiment of **Figure 4**, the old address of the SPVC endpoint destination node is identified in the Old Prefix field 404 and the new address of the SPVC endpoint destination node is identified in the New Prefix field 405.

[0038] In a further embodiment, the prefix fields 404, 405 are specified according to the NSAP node identification technique. The NSAP node identification technique, however, is often used to not only identify a particular node but also identify a particular port within a node. A port can be viewed as the architectural component of a node that collects traffic destined for a particular user or otherwise organizes the bandwidth of a node to finer degrees of granularity than the total bandwidth of the node.

[0039] For example (and referring briefly back to **Figure 1**), an SPVC having node 102₇ as its destination endpoint node may effectively set aside a portion of the bandwidth of link 103₁₄ for the use of the particular user associated with the SPVC. As such, an output port may be said to exist within node 102₇ that collects the information that is destined to be sent to the user over link 103₁₄ with this pre-defined bandwidth portion. Other output ports of node 102₇ may also be similarly identified for other SPVC users that are serviced by node 102₇ and consume other bandwidth portions of link 103₁₄.

[0040] Accordingly, if the NSAP identification technique is used to fill prefix fields 404 and 405, the prefix fields 404, 405 may not only identify the node that is experiencing an address change but may also identify some further granulated component of the node (e.g., such as an output port that services a particular

SPVC user). Nevertheless, the further granulized component can be effectively ignored by a node that receives the SIG information. That is, the network can automatically reconfigure itself based upon nodal information alone.

[0041] Accordingly, in various embodiments, a node that serves as a destination endpoint node for an SPVC can trigger the release of a PTSP having a PTSE with embedded SIG information that includes: 1) the previous address of the endpoint node; and 2) the new address of the endpoint node. In one embodiment, referring briefly back to Figure 1, the network management station 104 provides a change of address command to the destination endpoint node 102₇ (e.g., via an SNMP command or other technique). When the destination endpoint node 102₇ recognizes that its address has changed, it issues at least one PTSP to broadcast the fact that an address change is at hand.

[0042] **Figure 5** shows an embodiment of a methodology that may be executed by any of the other nodes 102₁ through 102₆ within the network 101 in response to the reception 501 of the PTSP. First, a receiving node attempts to determine whether or not it is configured to support an SPVC that is affected by the address change. In an embodiment, because SPVCs act as PVCs at the edges of an ATM network, only edge nodes attempt to make the above described determination.

[0043] That is, recalling that manual prior art network management efforts are directed to the endpoints of the SPVCs, these manual network management efforts may be eliminated if only those nodes within the network that could behave as an SPVC source endpoint actively use the SIG information of the

PTSP. As such, referring briefly back to **Figure 4**, in one embodiment the non edge nodes of the network (e.g., node 102₄ of **Figure 1**) are configured to ignore the SIG contents of a PTSP having nodal address change information once the ID_# field 403 of the SIG is recognized.

[0044] As seen in **Figure 5**, the nodal address found within the old prefix field 404 of the SIG information is compared 502 to the nodal address found within each of the SPVC prefixes currently supported by the node that received the PTSP. Where the nodal addresses match 503, an affected SPVC is found. That is, an SPVC is identified that: 1) is supported by the node receiving the PTSP; and 2) uses the node undergoing an address change as a destination endpoint node. For each match that occurs, the node that received the PTSP effectively replaces (within its SPVC records) the old nodal address information with the new nodal address information found within the new prefix field 405 of the SIG information.

[0045] Typically, a node is designed with a lookup table that lists the VPI/VCI information of each connection the node is currently configured to support. In an embodiment, the lookup table is configured to specify the destination endpoint address for each of those VPI/VCI listings that corresponds to an SPVC connection.

[0046] In this case, the aforementioned lookup table corresponds to the SPVC records referred to above; and, in order to implement the appropriate change, the lookup table will have the old destination endpoint address replaced with the new destination endpoint address. Once a node's SPVC records are updated,

subsequent SPVC connections will be properly routed to the appropriate destination endpoint by the node.

[0047] According to the PNNI approach, a node that receives any PTSE information can “re-issue” the PTSE information so that other nodes may receive it as well. If the lifetime of the PTSE information is not limited in some manner, the constant re-issuing of the PTSE will result in its never being removed from the network. Thus, various techniques may be employed to ensure that the PTSE is removed from the network (preferably sometime after, at least, each source endpoint node is able to receive it). In an embodiment, the lifetime of an issued PTSP message is deliberately limited by allowing each node that receives it to keep it in within it’s database for a limited time period. As is known in the art, the lifetime of PTSE information may be limited via manipulation of the PTSE lifetime field 311 originally shown in **Figure 3**.

[0048] Recall from above that a destination endpoint node can be configured to issue a PTSP having SIG information that describes its address change as soon as the destination endpoint node realizes its address is being changed. In another embodiment, the SIG information can be released with the next scheduled broadcast of PTSE information. That is, rather than release the SIG information upon an event (i.e., the address change), the SIG information is released as part of a scheduled (e.g., periodic) status update. In a further embodiment, the SIG information is embedded in the next scheduled broadcast of nodal PTSE information (e.g., that is periodically broadcast as part of the PNNI nodal status information sharing scheme).

[0049] Another aspect of the PNNI protocol is that it is easily scalable. That is, referring briefly back to **Figure 1**, the observed network 101 may actually be a small part of a much larger PNNI network that effectively interconnects various smaller networks (such as network 101) together. The interconnection as well as the other networks are not shown in **Figure 1** for convenience. According to the PNNI scheme, the smaller networks are referred to as peer networks.

[0050] As such, a larger PNNI network can be constructed by linking together a larger network of peer networks. Generally, a filtered or reduced flow of status/control information is shared between peers. That is, detailed status updating and event reporting (via the release of various PTSP packets) are exchanged within a peer network while less detailed status updating and event reporting are exchanged between peer networks. Nevertheless, in one embodiment, a PTSE having SIG information indicative of an edge node address change is exchanged between peer networks so that SPVCs that span across at least a pair of peer networks can be appropriately adjusted at an affected source endpoint node.

[0051] As routing and signaling protocols are often implemented with software, it is to be understood that embodiments of this invention may be used as or to support a software program executed upon some form of processing core (such as the CPU of a computer) or otherwise implemented or realized upon or within a machine readable medium. A machine readable medium includes any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer). For example, a machine readable medium includes

read only memory (ROM); random access memory (RAM); magnetic disk storage media; optical storage media; flash memory devices; electrical, optical, acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.); etc.

[0052] Furthermore, it is noteworthy to point out that a network node (which may also be referred to as a networking node, a node, a networking system and the like) is a system designed to act as a switch or a router or other device that relays information from a first networking line to a second networking line. A depiction of a networking node 600 is observed in **Figure 6**. A plurality of networking lines 601₁ through 601₆ (e.g., copper cables or fiber optic cables) are shown in **Figure 6** as being coupled to the networking node 600.

[0053] The node 600 is mostly responsible for collecting a traffic unit (e.g., a packet, a cell or a Time Division Multiplexed (TDM) time slot) from a first networking line (e.g., networking line 601₁) and re-transmitting at least a portion of it (e.g., its payload and various sections of its header) onto a second networking line (e.g., networking line 601₆). As such, the node 600 effectively relays information so that it may be carried over various geographic distances. Some degree of intelligence is involved in the relaying process so that the traffic units being collected are forwarded onto an appropriate networking line (e.g., in light of their source address and destination address).

[0054] As such, the node 600 of Figure 6 shows an traffic ingress/egress layer 602 and a switching/routing layer 603. The ingress/egress layer 602 is responsible for collecting inbound traffic units from the networking lines upon

which they arrived; and, presenting at least a portion of them (e.g., their header information) to the switching/routing layer 603. The ingress/egress layer 602 is also responsible for transmitting outgoing traffic units onto a networking line in response to the direction or control of the switching/routing layer 603.

[0055] The switching/routing layer 603 is responsible for effectively deciding which networking line is an appropriate networking line upon which a particular traffic unit should be transmitted upon. The switching/routing layer 603 often performs this activity based upon header information or other control information (such as SS7 based TDM connection information) associated with each traffic unit. Connection establishment and tear-down procedures (as well as network topology broadcasts or other networking overhead information) can often be viewed as being integrated into (or coupled to so as to communicate with) the switching/routing layer 603.

[0056] Note that the architecture of a networking system having a routing/switching layer 603 and an ingress/egress layer 602 may vary from embodiment to embodiment. For example, in some cases the switching/routing layer 603 may be designed onto a single card; or, in other cases, the switching/routing layer 603 may be designed across a plurality of cards. Also, in some cases the switching/routing layer 603 (or a portion thereof) may be integrated onto a Line Interface Card (LIC) that also acts as part of the ingress/egress layer 602.

[0057] In the foregoing specification, the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be

evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

039443-1270